

Amendments to the Claims:

This listing of claims will replace all prior versions, and listings, of claims in the application:

Listing of Claims:

1. (Currently amended) A method for distributing data among a plurality of data storage systems comprising:
 - producing profile information for a first data object that is stored in a first data storage system, said profile information comprising content-based information associated with said first data object;
 - communicating said profile information from the first data storage system to at least one second data storage system in said plurality of data storage systems;
 - calculating an interest metric at each of said at least one second data storage systems based on said profile information and on selection criteria maintained at said each of said at least one second data systems, wherein said selection criteria comprises information specific to said each of said at least one second data storage systems, said information including keywords received from a user at said each of said at least one second data storage systems;
 - receiving interest metrics at said first data storage system from said at least one of said second data storage systems;
 - producing an ordered set of target second data storage systems based on said interest metrics; and
 - communicating said first data object from the first data storage system to the first N of said target second data storage systems in said ordered set,
 - ~~communicating said interest metrics from each of said at least one second data storage systems to the first data storage system;~~
 - ~~selecting at least one target second data storage system at the first data storage system based upon the interest metrics; and~~
 - ~~copying said first data object to each target second data storage system;~~

wherein said first data object is copied to said each target second data storage system depending on content-based information associated with said first data object.

2. (Original) The method of claim 1 wherein said first data storage system comprises a server component in communication with a data storage component.

3. (Original) The method of claim 2 wherein said second data storage system comprises a server component in communication with a data storage component.

4-5. (Canceled)

6. (Previously presented) The method of claim 1, further comprising:
communicating said first data object to said target second data storage system if its interest metric exceeds a predetermined threshold.

7. (Previously presented) The method of claim 1, wherein said interest metric indicates whether or not to communicate said first data object to said target second data storage system.

8. (Previously presented) The method of claim 7 wherein if the interest metric indicates not to copy said first data object to said target second data storage system, then determining a replication site from among said second data storage systems independently of content of said first data object and copying said first data object to said replication site.

9. (Previously presented) The method of claim 18 wherein said selection criteria are stored in said first data storage system, said method further comprising communicating said first data object to said at least one target second data storage system based on said interest metric and a predetermined criterion.

10. (Original) The method of claim 9 further comprising additional selection criteria for an additional second data storage system, said method further comprising

communicating said first data object to said additional second data storage system based on said profile information and said additional selection criteria.

11. (Previously presented) The method of claim 18 wherein said selection criteria are stored in a selection server system separate from said first data storage system and from said second data storage system, said method further comprising:

communicating said profile information to said selection server system; and
receiving a selection indication from said selection server system,
wherein said first data object is selectively communicated to said second data storage system depending on said selection indication.

12. (Currently amended) A distributed data storage system comprising a plurality of data servers, each data server comprising:

a client interface component configured for communication with one or more clients to exchange data;

a data storage interface component configured for data communication with a data storage component; and

a data processing component configured to:

produce profile information associated with a first data object that is stored in said data storage component, said profile information comprising content-based information associated with content of said first data object;

communicate said profile information to a plurality of candidate data servers;

generate, at each of said plurality of candidate data servers, an interest metric based on the profile information and selection criteria maintained at each of said plurality of candidate data servers, wherein said selection criteria comprises information specific to said each of said plurality of candidate data servers, said information including keywords received from a user at said each of said plurality of candidate data servers;

receive interest metrics at said data storage component from each of said candidate data servers; and

select at least one candidate data server to receive said first data object based upon the interest metrics such that a candidate data server is selected if its corresponding interest metric exceeds a predetermined threshold; and

copy said first data object to the selected candidate data server.

~~copy said first data object to one or more of said candidate data servers based on the interest metrics received from said candidate data servers;~~

~~wherein an interest metric is produced by a candidate data server and is based on selection criteria stored in said candidate data server and on said profile information.~~

13-14. (Canceled)

15. (Previously presented) The data storage system of claim 12 wherein said interest metrics are binary indicators that indicate whether or not to copy said first data object to each of said candidate data servers.

16-17. (Canceled)

18. (Currently amended) A method for distributing data among a plurality of data storage systems comprising:

obtaining selection criteria in a first data storage system;

producing profile information for a first data object that is stored in said first data storage system, said profile information comprising content-based information associated with said first data object;

communicating the selection criteria and the profile information to at least one second data storage system;

generating, at each of said at least one second data storage systems, an interest metric based on the selection criteria and the profile information, wherein said selection criteria comprises information specific to said each of said at least one second data storage systems, said information including keywords received from a user at said each of said at least one second storage systems;

receiving ~~the~~ interest metrics at the first data storage system from said each of
said at least one second data storage systems;

producing an ordered set of target second data storage systems based on said
interest metrics; and

communicating said first data object to each of the first N of said target second
data storage systems in said ordered set if its corresponding interest metric exceeds a
predetermined threshold.

~~selecting at least one target second data storage system at the first data storage~~
~~system based upon the interest metric; and~~

~~copying said first data object to said at least one target second data storage~~
~~system.~~

19. (Original) The method of claim 18 further comprising receiving, at said
first data storage system, said selection criteria from one or more data storage systems other than
said first data storage system.

20. (Currently amended) A data system comprising:
a plurality of data centers; and
a plurality of client systems in data communication with said data centers,
each data center comprising:
a data storage component;
a file server component ~~operable~~operative to exchange data between a client
system and said data storage component;
a replicator component;
a receiver component; and
file selection criteria,
wherein said replicator component is ~~operable~~operative to produce profile data for
a data object that is to be replicated among one or more candidate target data centers, to
communicate said profile data to at least one of said candidate target data centers, to receive an

interest metric from each of said candidate target data centers, and to selectively communicate said data object to a candidate target data center ~~based on if~~ its interest metric exceeds a predetermined threshold, said profile data representative of content of said data object,

wherein said receiver component is ~~operable~~operative to receive the profile data information from a source data center and to ~~generate~~calculate an interest metric based on the profile data and selection criteria that is maintained in said receiver component and specific to the data center of which the receiver component is a part, said selection criteria comprising keywords received from a user of said data center, said receiver component further ~~operable~~operative to communicate the interest metric to said source data center for selectively copying said data object.

21. (Currently amended) The system of claim 20 wherein said candidate target data centers are ordered to produce an ordered set based on their corresponding interest metrics and said replicator component is further ~~operable~~operative to communicate said data object to the first N target data centers selected from said ordered set.

22. (Canceled)

23. (Previously presented) The system of claim 20 wherein said interest metric is an indication of whether or not to communicate said data object to said candidate target data center.

24. (Currently amended) A data system comprising:
a plurality of data centers; and
a plurality of client systems in data communication with said data centers,
each data center comprising:
a data storage component;
a file server component ~~operable~~operative to exchange data between a client system and said data storage component;
a replicator component; and

a collection of selection criteria comprising selection criteria provided from other data centers,

wherein said replicator component is ~~operable~~operative to produce profile data for a data object that is to be replicated among one or more candidate target data centers,

wherein said replicator component is operative to produce a plurality of interest metrics based on said collection of selection criteria and on said profile data, each interest metric corresponding to a data center in said plurality of data centers, and to produce an ordered set of said candidate target data centers based on their corresponding interest metrics,

wherein said replicator component is operative to produce an ordered set of said candidate target data centers based on their corresponding interest metrics, and to communicate said data object to each of the first N target data centers selected from said ordered set if its corresponding interest metric exceeds a predetermined value.

~~to communicate said profile data to at least one of said candidate target data centers, and to selectively communicate said data object to said candidate target data centers based on an interest metric corresponding to each of said candidate target data centers, said profile data representative of content of said data object, and~~

~~wherein at least one of said candidate target data centers is operable to receive the profile data, calculate the interest metric based on the profile data and said selection criteria, and communicate said interest metric to said replicator component.~~

25-28. (Canceled)

29. (Previously presented) A method for distributing data to a plurality of data storage systems in accordance with content-based interest metrics corresponding to each of said data storage systems, the method comprising:

producing a profile containing information representative of the content of a first data object stored in a first data storage system;

receiving interest information from a plurality of distinct second data storage systems specifying one or more categories of information requested for storage at each of said plurality of second data storage systems;

calculating interest metrics for each of the plurality of second data storage systems with respect to the first data object using the profile information and the interest information;

selecting one or more target second data storage systems to receive the first data object based upon their respectively calculated interest metrics; and

copying said first data object from said first data storage system to said one or more target second data storage systems.

30. (Previously presented) A method for distributing data among a plurality of data storage systems comprising:

receiving profile information representative of the content of a first data object stored in a first data storage system at a directory server;

receiving at the directory server interest information from a plurality of second data storage systems specifying one or more categories of information requested for storage at each of said plurality of second data storage systems;

calculating at the directory server interest metrics for each of the plurality of second data storage systems with respect to the first data object using the profile information and the interest information;

transmitting said interest metrics from the directory server to the first data storage system;

selecting at the first data processing system one or more target second data storage systems to receive the first data object based upon the interest metrics calculated at the directory server; and

copying said first data object to said one or more target second data storage system.

31. (Previously presented) The method of claim 1, further comprising updating metadata at the first data storage system with an identifier of the at least one target second data storage system to which the first data object is copied.

32. (Previously presented) The data storage system of claim 12 wherein said data processing component is configured to update metadata stored by said data storage component with identifiers of the one or more candidate data servers to which the first data object is copied.

33. (Previously presented) The method of claim 18 further comprising updating metadata at the first data storage system with an identifier of the at least one target second data storage system to which the first data object is copied.